arXiv:1402.5311v1 [cs.CC] 21 Feb 2014

On the Power of Multiplexing in Number-on-the-Forehead Protocols

Eldar Aharoni and Eyal Kushilevitz

Department of Computer Science, Technion, Haifa, Israel
{eldara|eyalk}@cs.technion.ac.il

Abstract. We study the direct-sum problem for k-party "Number On the Forehead" (NOF) deterministic communication complexity. We prove several positive results, showing that the complexity of computing a function f in this model, on ℓ instances, may be significantly cheaper than ℓ times the complexity of computing f on a single instance. Quite surprisingly, we show that this is the case for "most" (boolean, k-argument) functions. We then formalize two-types of sufficient conditions on a NOF protocol Q, for a single instance, each of which guarantees some communication complexity savings when appropriately extending Q to work on ℓ instances. One such condition refers to what each party needs to know about inputs of the other parties, and the other condition, additionally, refers to the communication pattern that the single-instance protocol Q uses. In both cases, the tool that we use is "multiplexing": we combine messages sent in parallel executions of protocols for a single instance, into a single message for the multi-instance (direct-sum) case, by xoring them with each other.

1 Introduction

Communication Complexity, introduced by Yao [25], studies the amount of communication parties need to exchange in order to carry out a distributed computation. In its most basic setting (so-called two-party communication complexity) two parties, Alice and Bob, each having n-bit input (x and y respectively) want to compute the value $f(x, y)$, for some fixed f. To do so, they use a predefined communication protocol Q. The cost of Q is the number of bits exchanged in Q, for the worst choice of (x, y). The communication complexity of a function f, denoted $D(f)$, is the minimal cost of an appropriate Q. The field deals with protocols for interesting functions and mostly with lower bounds for them, and has a wide range of applications (see [21] for an overview).

In the multi-party case, there are several common models. The "number on the forehead" (NOF) model of communication complexity, which is studied also in the current paper, was introduced by Chandra, Furst and Lipton [10]. In this model, k parties wish to jointly evaluate a function $f : (\{0,1\}^n)^k \to \{0,1\}$, where each input of f, i.e. $x_i \in \{0,1\}^n$ for $i \in [k]$, is given to all parties *except* P_i (pictorially, it is viewed as if x_i appears on the forehead of P_i). In order to compute f, the parties communicate by writing bits on a shared board. The communication complexity of a NOF protocol, is the number of bits written on the board, for the worst input. The NOF model is not as well understood (as, for example, the two-party case), mainly because of the complexity of dealing with the shared inputs, as every piece of the input is seen by all the parties but

one. An alternative multi-party model is the so-called "Number in Hand" model, where each x_i is given to party P_i (see, e.g., [13,23]). Even though the "Number in Hand" model may seem more natural, results for the NOF model have many applications in areas like branching programs complexity, boolean circuit complexity, proof complexity and pseudorandom generators. In all these applications, the overlap of information in the NOF setting, models well the problem in question. Note that for $k = 2$, both models collide with the standard two-party model. We next reference several papers, which studies NOF communication complexity. Separation of classes of complexity, within the NOF communication complexity, is examined, for example, in [5] and [11]. The NOF communication complexity of the disjointness function (defined in the following subsection) is studied in [6] [9] [22] and [24].

Direct-sum questions usually refer to the task of simultaneously solving multiple instances of the same problem, and ask whether this task can be done more efficiently than simply solving each instance separately. For example, the task of computing a function f on ℓ instances, is a direct-sum problem denoted f^ℓ. Direct-sum problems for two-party communication complexity have been first studied in [19], where the authors showed that if a certain direct-sum result holds for deterministic communication complexity of relations, then P $\neq$ NC^1. In [14], the authors proved that $D(f^\ell) \geq \ell \cdot (\sqrt{D(f)/2} - 2\log n - O(1))$ and presented a partial function f such that $D(f) = \Theta(\log n)$, whereas $D(f^\ell) = O(\ell + \log n \cdot \log \ell)$. They also proved that for any (even partial) function f, for one-round protocols, $D^1(f^\ell) \geq \ell \cdot (D^1(f) - O(\log n))$. In [20], the same is proved for the two-round deterministic case. Namely, in these cases no significant savings is possible. For general two-party deterministic protocols, this fundamental question (stated as a conjecture in [2]) remains open:

Question 1. Is it true that, for all functions $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, only a minor saving is possible; formally, $D(f^\ell) \geq \ell \cdot (D(f) - O(1))$?

While this is open for $k = 2$, we prove that the analogue question for the NOF model has negative answer, for every $k > 2$. Namely, savings is "usually" possible.

1.1 Our Results

Our first result shows that savings is possible for the direct-sum problem of "most" boolean functions, in the deterministic k-party NOF model. Namely, assume for convenience that $(k-1)$ divides ℓ, then $D^{\mathrm{NOF}}(f^\ell) \leq \ell \cdot \frac{n}{k-1} + \ell$, for all boolean functions f. Since "most" boolean functions have high communication complexity, specifically $D^{\mathrm{NOF}}(f) \geq n - \log(k+2)$ [15], then for most functions our upper bound is significantly cheaper than ℓ times the complexity of computing f on a single instance.

The above upper bound turns out not too difficult to prove. The main idea is that of appropriately *multiplexing* messages sent in various executions of the underlying protocol for f.

We proceed to further understanding the power of multiplexing. We formulate conditions under which multiplexing is possible. For this, we define a variant of the NOF model, in which parties can view only a subset of other parties' inputs. We give sufficient conditions on ℓ such protocols that guarantee that, instead of simply running

the ℓ protocol simultaneously, we can appropriately combine them into one (standard) NOF protocol that solves all ℓ instances, and this leads to some savings in communication. Next, we present sufficient conditions on a single protocol for f, which allow us extending it to a NOF protocol that solves ℓ instances of the original problem, with appropriate savings in communication complexity. The conditions used are technical, but we show, for instance, that for the important case of symmetric functions, they can be significantly simplified. We give examples of appropriate NOF protocols that are constructed by applying our technique.

We also use the so-called MYOPIC variant of the NOF model [16,8], which limits the way the inputs are shared between the parties and, additionally, limits the communication pattern between the parties to be one-way. We give conditions on ℓ MYOPIC protocols, which allow extending them to a NOF protocol solving ℓ instances, with appropriate savings in communication complexity.

1.2 Related Work

The direct-sum problem was also studied with respect to other communication complexity measures. For the non-deterministic case, it was shown [14] that $N(f^\ell) \geq \ell \cdot (N(f) - \log n - O(1))$. An alternative proof was given in [20]. For the randomized case, it is shown [14] that some savings is possible; specifically for EQ, the equality function, $R(EQ^\ell) = O(\ell + \log n)$, whereas $R(EQ) = \Theta(\log n)$. Using information-complexity arguments, Barak et al. [4] showed that (ignoring logarithmic factors) computing a function f on n instances, using a randomized protocol, requires $\Omega(\sqrt{n})$ times the randomized communication complexity of f.

Not much is known regarding the direct-sum problem f^ℓ, in the NOF model. In [6], the authors state that it is not clear how to prove a direct sum theorem for NOF. Certain direct-sum like operators, for the case of $n = 1$, were considered in [1]. The authors proved that for any symmetric boolean function $\mathrm{SYM} : \{0,1\}^\ell \to \{0,1\}$ (that depends only on the number of "1" inputs), and for any boolean function $f_k : \{0,1\}^k \to \{0,1\}$, the communication complexity $D(\mathrm{SYM}(f_k^\ell)) = O(\frac{\ell}{2^k} \cdot \log \ell + k \log \ell)$. The proof works for $k > 1 + \log \ell$. It is also shown that $D^{||}(\mathrm{SYM}(f_k^\ell)) = O((\log \ell)^3)$. Here, $D^{||}$ refers to the communication complexity in the simultaneous NOF model, where each party, without interaction, sends a single message to an external referee, who computes the output based only on the messages it received.

For randomized protocols, a strong direct product theorem for a problem, states that when bounding the communication complexity to be less than, say, ℓ times the communication complexity required for solving the problem on one instance, with constant error probability, then the probability of solving ℓ instances of the problem with constant error, is exponentially small in ℓ. Strong direct-product results in the two-party model, were shown for the one-way public-coin setting [17] and bounded round public-coin setting [18].

One of the most studied functions in communication complexity is the disjointness function, DISJ. Its k-argument variant DISJ_k is 1 iff, there is no $i \in [n]$ that belongs to all the k sets. That is, the inputs, viewed as sets, are disjoint. In one of the few results concerning direct-product in the NOF model, Sherstov [24] showed that $\ell \cdot \Omega(\frac{n}{4^k})^{\frac{1}{4}}$

bits are required to solve $(1-\epsilon)\cdot\ell$ out of ℓ instances of the function DISJ_k, for small enough $\epsilon > 0$. Clearly, solving all ℓ instances correctly, will take at least that much communication.

In [2,7], the communication complexity of other direct-sum like "operators" on multiple instances was studied. For example, the problem of choosing one instance from ℓ given instances and solving that instance. Alternatively, eliminating one ℓ-tuple of outputs which is *not* the correct one. In most of these cases, demonstrating savings turns out to be difficult.

2 Direct-Sum in the NOF Model – Basic Examples

As a first example, we show that savings is possible for the direct-sum problem of most boolean functions in the deterministic communication complexity NOF model. Specifically, for all functions whose NOF complexity is "high". From now on, we denote by $x_{i,j}$ the input on P_j's forehead in the i^{th} instance.

Lemma 1. *Let $f_k : (\{0,1\}^n)^k \to \{0,1\}$ be an arbitrary k-argument boolean function. Then, $D^{\mathrm{NOF}}(f_k^{k-1}) \leq n+k-1$.*

Proof. The idea is to use, what we term, "multiplexing". We present a protocol Q for the problem: P_k computes $M = \oplus_{i=1}^{k-1}(x_{i,i})$. That is, it xores the input of P_1 for instance 1 with the the input of P_2 for instance 2, and so on for each of the $k-1$ instances. Party P_k writes this value on the common board. Now, for $1 \leq i < k$, party P_i knows all the values used by P_k to compute M, except for $x_{i,i}$ so, by xoring M with those $k-2$ values, P_i can determine $x_{i,i}$. This, in turn, means that now P_i knows all the inputs for instance i, and so it locally computes the output of f on the i-th instance and writes this value on the board. As f's value on each instance, has been written by the corresponding P_i ($1 \leq i < k$), then all parties now know the correct output for each of the $k-1$ instances. □

If the number of instances, ℓ, is different than $k-1$, we simply partition the instances into "blocks" of size $k-1$. Hence,

Corollary 1. *Let $f_k : (\{0,1\}^n)^k \to \{0,1\}$ be an arbitrary k-party boolean function and $(k-1)$ divides ℓ, then $D^{\mathrm{NOF}}(f_k^\ell) \leq \ell \cdot \frac{n}{k-1} + \ell$.*

In [15], it is proved that most k-argument functions $f : (\{0,1\}^n)^k \to \{0,1\}$, have communication complexity $D^{\mathrm{NOF}}(f) \geq n - \log(k+2)$, in the NOF model (see also [21, Exercise 6.5]). Using Corollary 1 with such a function, gives a provable savings for the direct-sum problem in the deterministic NOF model. Since most functions are also hard with respect to other complexity measures in the NOF model (e.g., randomized complexity, as showen in [12]), Corollary 1 implies a similar result for those as well.

The above example deals with savings for "hard" functions. Next, we show how multiplexing can (slightly) help in cases where non-trivial protocols exist for a single instance. Specifically, we consider the boolean function "equality", which is defined as $EQ_k(x_1,\ldots,x_k) = 1$ iff $x_1 = \ldots = x_k$.

Proposition 1. $D^{\mathrm{NOF}}(EQ_k) = 2$.

Proof. We present a protocol for the problem. Party P_{k-1} writes on the board 1 if $x_{k-2} = x_k$, and 0 otherwise. If P_k sees 1 on the board and if all the inputs it sees are equal, then P_k writes as output "1", otherwise it writes "0". Obviously, the protocol is correct, so $D^{\mathrm{NOF}}(EQ_k) \leq 2$. For the lower bound, the first party P_i to write a bit on the board (the index of this party is fixed by the protocol), must sometimes also read something from the board to know the value of EQ_k (in case that all other inputs x_j are equal), so that requires at least two bits of communication. □

Proposition 2. *For odd $k \geq 3$, we have $D^{\mathrm{NOF}}(EQ_k^{\frac{k-1}{2}}) \leq 1 + \frac{k-1}{2}$.*

Proof. We give a protocol for $EQ_k^{\frac{k-1}{2}}$: For odd j's, define

$$b_j = \begin{cases} 1 \text{ if } x_{j,j} = x_{j,j+1} \\ 0 \text{ otherwise} \end{cases}$$

- P_k computes the xor of all the bits b_j, for $j = \{1, 3, \ldots, k-2\}$, which we denote as M, and writes in on the board.
- For $j \in [1, \frac{k-1}{2}]$, party $P_{2 \cdot j}$ knows all the values used to compute M except $b_{2 \cdot j-1}$, so it can now compute it. Then, if all the inputs in instance $2 \cdot j$ that $P_{2 \cdot j}$ sees are the same and $b_{2 \cdot j-1} = 1$, then it writes 1 on the board as the output for this instance; otherwise, it writes 0.

The communication is $1 + \frac{k-1}{2}$ bits (i.e., better than $\frac{k-1}{2} \cdot D^{\mathrm{NOF}}(EQ_k) = k - 1$.) □

3 Multiplexing NOF protocols via a Restricted NOF Model

The above examples already show the power of multiplexing. In this section, we will describe how to multiplex the executions of a single-instance NOF protocol Q on ℓ instances, assuming that Q meets certain conditions that would allow it to run in a "restricted NOF" model. Before formally defining the restricted NOF model and how it is used, we describe, informally, the basic ideas.

As a starting point, consider a 3-party NOF protocol Q, for a boolean function $f \in (\{0,1\}^n)^3 \to \{0,1\}$. Assume that the first message of Q is written on the board by P_1, and that it is of the form $m_2(x_2)||m_3(x_3)$, where $||$ stands for concatenation. Assume also that m_2, m_3 are of the same length (for all inputs). That is, P_1's message has two equal-length parts, one that depends on input x_2 and one that depends on input x_3. Obviously, P_2 can compute by itself the message $m_3(x_3)$, as it sees x_3, and, similarly, P_3 can compute $m_2(x_2)$. Therefore, in this specific case, instead of writing both these messages, P_1 could have simply written on the board $m_2(x_2) \oplus m_3(x_3)$ (which is of half the size), and both P_2 and P_3 could have determined the original message they were supposed to get. This can be extended beyond the first message of the protocol: namely, when m_2 depends on x_2 and on the communication so far, and m_3 depends on

x_3 and on the communication so far. Since both P_2 and P_3 see the board themselves, they can still compute the values m_3, m_2, respectively.

Next, we wish to extend the above simple example to more parties; specifically, consider the case $k = 4$. Suppose that P_1 would like to write on the board the expression $m_2(x_2, x_4)||m_3(x_3, x_4)$. In this case, P_2 and P_3 will still be able to determine m_2, m_3 even if P_1 will write on the board the value $m_2 \oplus m_3$. However, P_4 (who does not know x_4) cannot compute m_2, m_3 from m. In some cases, however, P_4 itself does not need the values $m_2.m_3$ to continue its computation. To model what information is needed by each party, it is convenient to view the NOF protocol Q as a protocol Q' in a model where communication is over point-to-point channels (where a party that does not need a message, simply does not receive it). Say that in Q' party P_1 sends a message $m_2(x_2, x_4)$ to P_2 and a message $m_3(x_3, x_4)$ to P_3. We then can get a new NOF protocol Q'' by multiplexing this message, and so we save the length of a single message. In Q'', party P_1 will write on the board $M = m_2(x_2, x_4) \oplus m_3(x_3, x_4)$. Now, P_2 and P_3 will be able to determine their original messages, as explained before, and since we assumed that in Q', party P_4 did not receive a message at this stage, its future messages in Q' do not depend on messages m_2, m_3 and so P_4 can simply ignore the value M on the board in Q''.

Now that we have moved to the point-to-point model, we need to reconsider what if the messages we wanted to multiplex was not the first message to be sent. For instance, suppose we take a protocol which begins with P_2 and P_3 sending a message, each, to P_1. Then, P_1 sends a message $m_2(x_2, H)$ to P_2 and a message $m_3(x_3, H)$ to P_3, where H stands for the history of messages that P_1 received. Now, if P_1 will multiplex these messages, then P_2 may not be able to compute $m_3(x_3, H)$ as it may have only partial information on H (since we are dealing now with a point-to-point model, P_2 did not see the message sent from P_3 to P_1). However, when having, say, two instances for f, we can sometimes still do something. Below, we formalize constrains for when two (equal length) messages m_1 in protocol Q_1 and m_2 in Q_2 can be multiplexed. Obviously, the sender of these two messages should be the same, say P_i, and the recipients need to be different, say P_{j_1} and P_{j_2}. Denote these messages by $m_{1,i\to j_1}$ and $m_{2,i\to j_2}$. In order for P_{j_1} to be able to compute the message $m_{2,i\to j_2}$ (in order to "de-multiplex" the message on the board), it is enough that (1) this message does not depend on x_{j_1}, and (2) that P_{j_1} knows the history of messages received by P_i so far (similar requirements needed for P_{j_2}). To fulfil the second requirement, we require that none of the messages sent to P_i in Q_1 and Q_2 will be multiplexed (so in the resulting NOF protocol, those messages will simply appear on the board and P_{j_1} will know them). This discussion can be simply extended to deal with ℓ protocols for which we want to multiplex messages from P_i to $P_{j_1}, \ldots, P_{j_\ell}$, respectively.

To force protocols to satisfy the above requirements, we define below a new model, where instead of having requirements on messages sent from parties, we simply limit the parties' view of other parties' inputs. This is defined using a directed graph G, that describes what inputs are seen by each party. Still, communication is allowed between any pair of parties over a private channel.

Definition 1. *Let $G = (V, E)$ be a simple directed graph over $V = \{1, \ldots, k\}$. A restricted NOF_G multi-party model with k parties has the restriction that, for all $i, j \in$*

$[k]$, party P_i sees the input of party P_j iff $(i,j) \in E$. The graph G will be referred to as the restriction graph. Communication in the NOF_G model is over a complete network of point-to-point channels. Namely, in a NOF_G protocol Q, in each round t, each party P_i can send a message to each party P_j depending on messages P_i received in the previous $t-1$ rounds and depending on the inputs P_i sees. (We emphasize that we allow P_i sending a message to P_j even though $(i,j) \notin E$). We say that protocol Q, in the NOF_G model, computes $f \in (\{0,1\}^n)^k \to \{0,1\}$, if some party P_i, whose identity is known to all parties, outputs $f(\boldsymbol{x})$.[1]

We will use the following notations. For a node $i \in V$ denote the neighbors of i, by $N(i) = \{j|(i,j) \in E\}$. For an input $\boldsymbol{x}$, denote $\mathrm{VIEW}_{i,G}(\boldsymbol{x})$ as the inputs that P_i sees, that is, the sequence $\{x_j|j \in N(i)\}$. Denote $m_{j\to i,t,\boldsymbol{x}}$ as the message sent from P_j to P_i in round t on input $\boldsymbol{x}$. Let $M_{i,t}(\boldsymbol{x}) = \{(t', j, m_{j\to i,t',\boldsymbol{x}})|t' \le t, j \in [k]\}$, be the history of messages that P_i received in rounds 1 to t on input $\boldsymbol{x}$. The message sent from P_i to P_j in round $t+1$ in protocol Q may depend on the messages P_i received in previous rounds and on the inputs seen by P_i (as set by G). That is, it can be denoted as $m^Q_{i\to j,t+1,\boldsymbol{x}}(M_{i,t}(\boldsymbol{x}), \mathrm{VIEW}_{i,G}(\boldsymbol{x}))$. When the input $\boldsymbol{x}$ is clear from the context, we will usually omit it from the notations.

Definition 2. *An oblivious k-party NOF_G protocol Q is a protocol in which, for every input vector $\boldsymbol{x}$, the length of the messages exchanged is fixed. That is, there exist a mapping $\mathrm{LEN}^Q : \mathbb{N} \times [k] \times [k] \to \mathbb{N}$, referred to as the* communication pattern of Q*, that maps (t,i,j) to the length of the message sent from P_i to P_j in round t, on every $\boldsymbol{x}$. For an oblivious protocol Q, we denote the number of bits sent from P_a to P_b in protocol Q by $W_Q(a,b)$.*

The following definitions will capture sufficient requirements from restricted NOF protocols, so that it is possible to multiplex messages when combining the protocols together.

Definition 3. *Let $G = (V,E)$ be a restriction graph for k parties. Denote the isomorphic to G graph, which is the resulted graph of permutating G's nodes with a permutation π, by G_π. Let $\overline{N_G(a)}$ be the non-neighbours of a in the graph G, that is, $V \backslash N_G(a)$. Let $\prod^\ell = \pi_1, \ldots, \pi_\ell$ be a sequence of ℓ permutations over $[k]$. For $b \in [k], R \subseteq [\ell]$, denote $\mathrm{MAP}_{\prod^\ell,b,R} = \{\pi_r(b) : r \in R\}$. That is, the images of b under the permutations in R. A triplet (a,b,R), for $a \ne b \in [k]$ and $R \subseteq [\ell]$, is $\prod^\ell_G$-good if:*

1. $\mathrm{MAP}_{\prod^\ell,a,R} = \{a\}$.
2. $\mathrm{MAP}_{\prod^\ell,b,R}$ *is a subset of* $\overline{N(a)}$ *of size* $|R|$*, that does not include* a *or* b*.* [2]
3. *For all* $r \in R$ *we have* $(\mathrm{MAP}_{\prod^\ell,b,R} \bigcup \{b\} \backslash \{\pi_r(b)\}) \subseteq \overline{N_{G_{\pi_r}}(a)}$.

Let us discuss the above definition. Suppose we have a NOF_G protocol Q that we want to multiplex. Intuitively, each triplet (a,b,R) is a candidate for multiplexing. The

[1] This is in contrast to the NOF model, where all the parties output $f(\boldsymbol{x})$.

[2] Condition (1) already implies that $a \notin \mathrm{MAP}_{\prod^\ell,b,R}$, as a is already mapped to a by all the permutations in R. This requirement is included in condition (2) for convenience.

messages that will be multiplexed, are the message that P_a sends to P_b in Q, and the messages $P_{\pi_r(a)}$ sends to $P_{\pi_r(b)}$, for each permutation π_r, for $r \in R$. Condition (1) guarantees that a is mapped to itself in all those permutations, so it is the only sender of the messages we want to multiplex together, as required for multiplexing. Condition (2) requires that $b \notin \mathrm{MAP}_{\prod^\ell,b,R}$ and $\mathrm{MAP}_{\prod^\ell,b,R}$ being the size of $|R|$, which guarantees that we will have $|R|+1$ messages to multiplex (messages which are sent from P_a to P_b, or to other "alternative" recipients, as determined by R). As stated before, for multiplexing to work, any multiplexed message sent to a certain recipient, should not depend on any of the other recipients' input. Conditions (2) (i.e., that $\mathrm{MAP}_{\prod^\ell,b,R}$ is a subset of $\overline{N(a)}$) and condition (3) take care of that.

Definition 4. *Let G, k, and $\prod^\ell$ be as in Definition 3. A $\prod_G^\ell$-multiplexing set S is a collection of $\prod_G^\ell$-good triplets, such that for any two distinct $(a,b,R),(a',b',R') \in S$ we have $a \notin (\{b'\} \bigcup \mathrm{MAP}_{\prod^\ell,b',R'})$ and $(\{b\} \bigcup \mathrm{MAP}_{\prod^\ell,b,R}) \bigcap (\{b'\} \bigcup \mathrm{MAP}_{\prod^\ell,b',R'}) = \emptyset$.*

Intuitively, a $\prod^\ell$-multiplexing set S guarantees that if we multiplex any messages sent from a, we will not multiplex messages sent *to* a. Also, we will not have a recipient or an "alternative" recipient, for the same a, more than once.

The next definition formalizes the "communication pattern" of an oblivious protocol. Having the same "communication pattern" between protocols is necessary for our multiplexing stage in the following proofs.

Definition 5. *Let Q be an oblivious k-party NOF_G protocol (for some restriction graph $G=(V,E)$), which computes a k-argument boolean function f. For a communication pattern LEN, we denote LEN_π as the mapping $\mathrm{LEN}_\pi(t,i,j) = \mathrm{LEN}(t,\pi^{-1}(i),\pi^{-1}(j))$. Protocol Q is π-pattern-robust, if there is an oblivious protocol Q' for computing f in NOF_{G_π}, such that $\mathrm{LEN}^{Q'} = \mathrm{LEN}_\pi^Q$. That is, if there is a NOF_{G_π} protocol Q' for computing f, such that the communication pattern is the same for the two protocols, up to the permutation π.*

Theorem 1. *Let f be a k-argument boolean function, $G=(V,E)$ a restriction graph for k parties, Q an oblivious k-party NOF_G protocol for f and $\prod^\ell = \pi_1,\ldots\pi_\ell$ be a sequence of permutations over $[k]$ such that π_1 is the identity permutation. Also, assume that, for $2 \le i \le \ell$, protocol Q is π_i-pattern-robust, and S is a $\prod_G^\ell$-multiplexing set. Then,*

$$D^{\mathrm{NOF}}(f^\ell) \le \ell \cdot D^{\mathrm{NOF}_G}(Q) + \ell - \sum_{(a,b,R)\in S} |R| \cdot W_Q(a,b).$$

Intuitively, the NOF communication complexity of f^ℓ is bounded by ℓ times the NOF_G communication complexity of f, plus ℓ bits to distribute the outputs, minus the savings, which depends on the multiplexing set. For any (a,b,R) in the multiplexing set, we save $|R|$ times the communication between a and b in the protocol Q.

Proof. Protocol Q is π_i-pattern-robust, for $i=2,\ldots,\ell$. Therefore, there are protocols Q_{π_i} such that Q_{π_i} solves f in $\mathrm{NOF}_{G_{\pi_i}}$ and $\mathrm{LEN}^Q = \mathrm{LEN}^{Q_{\pi_i}}$. We take those ℓ protocols and combine them into a single NOF protocol Q' for computing f^ℓ, where the

i^{th} protocol is responsible for computing f on the i^{th} instance. Whichever messages are sent in $Q_{\pi_1}, \ldots, Q_{\pi_\ell}$ over private channels are now communicated on the board, so clearly the computation of each message remains the same (as the messages previously received by a party, are now known to all). According to this definition of Q', for each $(a_i, b_i, R_i) \in S$, the bits initially sent in round t of protocol Q from P_{a_i} to P_{b_i}, are now translated into $|R_i| + 1$ messages. That is, one message from P_{a_i} to P_{b_i} in Q itself, and one message from P_{a_i} to $P_{\pi_r(b_i)}$ in Q_{π_r}, for $r \in R_i$. Instead of sending all these messages in Q', we perform the following multiplexing step for every round t: Party P_{a_i} writes on the board

$$m^{Q'}_{(a_i,b_i,R_i),t} \triangleq m^{Q}_{a_i \to b_i,t} \oplus \bigoplus_{r \in R_i} (m^{Q_{\pi_r}}_{a_i \to \pi_r(b_i),t}).$$

We need to prove that (i) for each $r \in R_i$, party $P_{\pi_r(b_i)}$ can infer the original message sent to it in Q_{π_r}, and that (ii) P_{b_i} itself can retrieve the message originally sent to it in Q. We start with (i). To prove that $P_{\pi_r(b_i)}$ can compute the message $m^{Q'}_{a_i \to \pi_r(b_i),t}$, we need to show that:

1. All the xored messages (beside $m^{Q_{\pi_r}}_{a_i \to \pi_r(b_i),t}$) do not depend on the input of $P_{\pi_r(b_i)}$ (note that since Q' is in the regular NOF model, $P_{\pi_r(b_i)}$ knows all the inputs but its own).
2. $P_{\pi_r(b_i)}$ knows the history of messages received by P_{a_i} in all those protocols (as the xored messages may also depend on them).

For (1), let $r \neq r' \in R$. We show that $P_{\pi_r(b_i)}$ can compute $m^{Q_{\pi_{r'}}}_{a_i \to \pi_{r'}(b_i),t}$. This follows from the fact that $\pi_r(b_i)$ belongs to $\mathrm{MAP}_{\prod^\ell, b_i, R_i} \bigcup \{b_i\} \backslash \{\pi'_r(b_i)\}$ (by definition of $\mathrm{MAP}_{\prod^\ell, b_i, R_i}$), and $(\mathrm{MAP}_{\prod^\ell, b_i, R_i} \bigcup \{b_i\} \backslash \{\pi'_r(b_i)\}) \subseteq \overline{N_{G_{\pi'_r}}(a_i)}$ by the $\prod^\ell_G$-goodness of triplet (a_i, b_i, R_i)) (specifically, Definition 3, condition (4)). We show that $P_{\pi_r(b_i)}$ can compute $m^{Q}_{a_i \to b_i,t}$. This follows from the fact that $\mathrm{MAP}_{\prod^\ell, b_i, R_i} \subseteq \overline{N(a_i)}$ (by Definition 3, condition (2)). For (2), this is true, as messages which are sent to P_{a_i} in the original protocols, are not multiplexed in the combined protocol (by Definition 3, for $u \neq i$, we have that $a_i \notin (\{b_u\} \bigcup \mathrm{MAP}_{\prod^\ell, b_u, R_u})$.

The proof for (ii) is similar, but somewhat simpler. The argument for (2) is just the same. The argument for (1), is as follows: Let $r \neq r' \in R$. We show that party P_{b_i} can compute $m^{Q}_{a_i \to \pi_{r'}(b_i),t}$. This is true because b_i belongs to $\mathrm{MAP}_{\prod^\ell, b_i, R_i} \bigcup \{b_i\} \backslash \{\pi_{r'}(b_i)\}$ (as $b_i \neq \pi_{r'}(b_i)$, by Definition 3, condition (2)).

An additional step in Q' is the reporting of the output for each of the instances, to all parties. As the protocols Q and Q_{π_i}, for $i = 2, \ldots, \ell$, all ended up with one party knowing the output, and since Q' is in the NOF model, then for each $i = 1, \ldots, \ell$, the party who knows the output on instance i now writes the output bit on the board. It follows that $D^{\mathrm{NOF}}(f^\ell) \leq \ell \cdot D^{\mathrm{NOF}_G}(Q) + \ell - \sum_{(a,b,R) \in S} |R| \cdot W_Q(a, b)$. □

Example 1. We now show that Theorem 1 implies Lemma 1. Let f be a k-argument function. Let $\ell = k - 1$. Let $G = ([k], \{(k, 1)\} \bigcup \{(1, i) : 2 \leq i \leq k\})$ be a restriction graph. I.e, the only edges in G are between k to 1 and between 1 to all the other nodes. We describe a protocol Q for f in NOF_G: P_k sends x_1 to P_1 and P_1 computes $f(\boldsymbol{x})$,

as it now knows all k inputs. We take permutations $\pi_2 = (2, 3, \ldots, k-1, 1, k), \pi_3 = (3, 4, \ldots, k-1, 1, 2, k), \ldots, \pi_\ell = (k-1, 1, \ldots, k-2, k)$. Let Q^i be the following protocol for computing f in $\mathrm{NOF}_{G_{\pi_i}}$: party P_k sends P_i its input, and P_i (which now knows all the inputs) computes the value $f(\boldsymbol{x})$ and outputs it. As $\mathrm{LEN}^{Q^i} = \mathrm{LEN}^Q_{\pi_i}$, it follows that Q is π_i-pattern-robust. Let S contain one triplet $(k, 1, \{2, \ldots, k-1\})$. We argue that the triplet is $\prod_G^\ell$-good. The value k is mapped to itself in all the permutations. The value 1 is mapped to $k-2$ different values under the $k-2$ permutations $\pi_2, \ldots, \pi_l$, none of which is 1 and all of which are not neighbours of k in G. We have $\mathrm{MAP}_{\prod^\ell,1,R} = \{2, \ldots, k-1\}$, and for all $r \in \{2, \ldots, k-1\}$, we have $N_{G_{\pi_r}}(k) = \{r\}$. The condition that for all $r \in \{2, \ldots, k-1\}$, $(\mathrm{MAP}_{\prod^\ell,1,R} \bigcup \{1\} \backslash \{\pi_r(1)\}) \subseteq \overline{N_{G_{\pi_r}}(k)}$, translates to requiring that for for all $r \in \{2, \ldots, k-1\} : [k-1] \backslash \{r\} \subseteq [k-1] \backslash \{r\}$, which is true. Then, by Theorem 1, it follows that $D^{\mathrm{NOF}}(f_k^\ell) \leq \ell \cdot n + \ell - (k-2) \cdot n = n + \ell$.

Employing Theorem 1, requires having ℓ protocols for f in ℓ restricted NOF models. Sometimes finding such protocols may be difficult. We next turn to a specific family of functions, specifically the symmetric functions, and see how we can have a simpler theorem that can be applied to functions in this family.

Definition 6. *A π-symmetric k-argument function $f \in (\{0,1\}^n)^k \to \{0,1\}$ is a function whose value remains the same when applying the permutation π to its inputs. That is, for any $\boldsymbol{x} \in (\{0,1\}^n)^k$, it holds that $f(x_1, \ldots x_k) = f(x_{\pi(1)}, \ldots, x_{\pi(k)})$.*

Definition 7. *A symmetric k-argument function $f \in (\{0,1\}^n)^k \to \{0,1\}$ is a function who that is π-symmetric, for all permutations π over $[k]$. That is, a function which does not depend on the order of its inputs.*

Let f be a k-argument function and Q a NOF protocol for f. Imagine each party P_i has its name changed to $P'_{\pi(i)}$. Now, suppose that given input $\boldsymbol{x}$, the parties simply execute Q, only using their new identities, instead of the old ones. For input vector $\boldsymbol{x}$, protocol Q' will correctly compute $f(\pi^{-1}(\boldsymbol{x}))$. Now, view the original Q, as a NOF_G protocol for a some restriction graph G. The same manipulation can be applied to protocol Q, to get a NOF_{G_π} protocol Q'. This is true, as $\mathrm{VIEW}_{i,G_\pi}(\boldsymbol{x}) = \mathrm{VIEW}_{\pi^{-1}(i),G}(\pi^{-1}(\boldsymbol{x}))$. We formalize this below.

Definition 8. *Let f be a k-argument function, Q be a k-party NOF_G protocol for a given graph $G = (V, E)$ and π a permutation over $[k]$. Denote $\pi(Q)$ as a protocol based on Q for a restricted NOF_{G_π}, such that for all i, j, t the message $m^{\pi(Q)}_{i \to j,t}$ for the input $\boldsymbol{x}$ is the corresponding $m^Q_{\pi^{-1}(i) \to \pi^{-1}(j),t}$ for the input $\pi^{-1}(\boldsymbol{x})$.*

To show that this protocol is well defined, we argue that the function $m^{\pi(Q)}_{t+1,i \to j}$ is computable by P_i in NOF_G. For that, we need to show that P_i at round $t+1$ knows $M_{\pi^{-1}(i),t}$ and $\mathrm{VIEW}_{\pi^{-1}(i),G}$. This can be shown by a simple induction, where the base step ($t = 1$) is correct as $\mathrm{VIEW}_{i,G_\pi}(\boldsymbol{x}) = \mathrm{VIEW}_{\pi^{-1}(i),G}(\pi^{-1}(\boldsymbol{x}))$ and $M_{\pi^{-1}(i),0} = \emptyset$. This implies the following simple claim (proof omitted).

Claim 1. Let f, k, Q, G, π be as in Definition 8. If f is π-symmetric, then $\pi(Q)$ is a correct NOF_{G_π} protocol for f.

Definition 9. *Let $G = (V, E)$ be a restriction graph for k parties. A set of triplets*

$$S \subseteq \{(a, b, B) | a \neq b \in [k], B \subseteq \overline{N}(a), a, b \notin B\}$$

is a filtering set for G, if for any two distinct $(a, b, B), (a', b', B') \in S$, we have that $a \notin (\{b'\} \bigcup B')$ and $(\{b\} \bigcup B) \bigcap (\{b'\} \bigcup B') = \emptyset$. For a filtering set S, denote $R(S) = \max_{a \in [k]} \sum_{(a,b',B') \in S} |B'|$. A filtering set for G is an ℓ-filtering set, if $R(S) \leq \ell - 1$.

Let f be a k-argument function and $G = (V, E)$ be a restriction graph for k parties. We now have two, somewhat similar, definitions: an ℓ-filtering set for G and a $\prod_G^\ell$-multiplexing set. We discuss the difference between them. Finding a $\prod_G^\ell$-multiplexing set may be difficult, as the definition requires us to have a sequence of ℓ permutations. Furthermore, we need to have ℓ oblivious $\mathrm{NOF}_{G_{\pi_r}}$ protocols which computes f, a protocol for each permutation π_r, such that they all have the same LEN (up to the permutation). The multiplexing set, is a set of potential triplets, representing which protocols and which messages in those protocols, we can multiplex. Instead, for the special case of symmetric functions, to get a multiplexed protocol, we only need to have one oblivious NOF_G protocol. The filtering set will have triplets, representing multiplexing candidates, based only on the graph G, and the permutations themselves will be computed from the filtering set.

Theorem 2. *Let f be a k-argument symmetric function, $G = (V, E)$ a restriction graph for k parties, Q an oblivious NOF_G protocol for f and S an ℓ-filtering set for G. Then,*

$$D^{\mathrm{NOF}}(f^\ell) \leq \ell \cdot D^{\mathrm{NOF}_G}(Q) + \ell - \sum_{(a,b,B) \in S} |B| \cdot W_Q(a, b).$$

Proof. We start by presenting an overview of the proof. Our aim is to form a sequence of ℓ permutations π_r over $[k]$, written as an $\ell \times k$ matrix, where each row is a permutation. We then take those permutations and find for each permutation π_r, a $\mathrm{NOF}_{G_{\pi_r}}$ protocol Q_{π_r} which is π_r-pattern-robust with respect to Q, and which solves f. We then convert the given ℓ-filtering set S into a respective $\prod_G^\ell$-multiplexing set S^*, for the aforementioned permutations and apply Theorem 1 to conclude the proof.

We construct an $\ell \times k$ matrix, $A_{G,\ell,S}$ (or A in short), such that the first row of A is the identity permutation over $[k]$. Let $S = \{(a_1, b_1, B_1), \ldots, (a_{|S|}, b_{|S|}, B_{|S|})\}$ be the given ℓ-filtering set. Each triplet (a_i, b_i, B_i), will partially fill entries in A, so that in the corresponding permutations, b_i is mapped to elements of B_i. Specifically, denote $last_i = 1 + \sum_{u<i | a_i = a_u} |B_u|$, that is, the last index of a row in A, which we filled with entries having to do with a_i. Assume $a_i = a_u$ and, without loss of generality, $i < u$, then for $j = 1, \ldots, |B_i|$ its follows from the definition of $last$ that $last(i) + j \leq last(u)$. For $j = 1, \ldots, |B_i|$, denote $row(i, j) = last_i + j$. For the triplet (a_i, b_i, B_i), for $1 \leq j \leq |B_i|$, entries will be added to rows $row(i, j)$. Eventually, $row(i, j)$, will hold a permutation, where a_i will be mapped to itself and b_i will be mapped to $(B_i)_j$. For example, if all $a_i's$ are different, then $last_i = 1$ for all i. If, for example, all the $a_i's$ are the same, each $row(i, j)$ is a different row in the matrix. See Example 2 below for a set S, its corresponding matrix and the corresponding function row.

More formally, for $(a_i, b_i, B_i) \in S$, for $j = 1, \ldots, |B_i|$ we have $A[row(i,j)][a_i] = a_i$ and $A[row(i,j)][b_i] = (B_i)_j$. Note that since $R(S) \leq \ell - 1$, we have that $row(i,j) \leq l$, and so, ℓ rows suffice for A. For a columns set $T \subseteq [k]$, we will use the notation $A[i][T]$ for $\{A[i][j] : j \in T\}$. We claim that the above A can be extended to accommodate the following restrictions:

1. Each row of A is a permutation over $[k]$.
2. For each $1 \leq i \leq |S|, 1 \leq j \leq |B_i| : A[row(i,j)][B_i] = B_i \cup \{b_i\} \setminus \{(B_i)_j\}$.

First we argue that so far, there is no row in A with a value appearing more than once. Since the sets $\{b_i\} \bigcup B_i$ are always disjoint (from one another, and from $\{a_i\}$'s, as S is a filtering set), the only possible repetition can be between $a_i's$. That is, there are $i_1 < i_2, j_1, j_2$ such that $1 \leq j_1 \leq |B_{i_1}|, 1 \leq j_2 \leq |B_{i_2}|, a_{i_1} = a_{i_2}$ and $row(i_1, j_1) = row(i_2, j_2)$. That is, $last(i_1) + j_1 = last(i_2) + j_2$. As explained before, $last(i_1) + j_1 \leq last(i_2) < last(i_2) + j_2$, so that leads which is a contradiction. Given that currently A does not have any repeating value in any of its rows, and that the above two requirements of the extension, only pose restrictions per row, on values which currently do not appear in the row, it is obvious that A can be completed as required.

From now on, we can assume that for each $1 \leq r \leq \ell$, the row $A[r]$ is a permutation which will be denoted as π_r and the protocol Q_{π_r} is a suitable $\mathrm{NOF}_{G_{\pi_r}}$ protocol for computing f, as guaranteed by Claim 1. That is, Q itself is π_r-pattern-robust.

For $(a_i, b_i, B_i) \in S$, let $R_i = \{row(i,j) : 1 \leq j \leq |B_i|\}$ and $S^* = \{(a_i, b_i, R_i) : (a_i, b_i, B_i) \in S\}$. We prove that S^* is a $\prod_G^\ell$-multiplexing set. We start by proving that for every $1 \leq i \leq |S|$, the triplet (a_i, b_i, R_i) is $\prod_G^\ell$-good.

Condition (1) of Definition 3 is met, as A was built such that $A[row(i,j)][a_i] = a_i$. $\mathrm{MAP}_{\prod^\ell, b_i, R_i}$ is simply $\{\pi_r(b_i) : r \in R_i\} = \{\pi_r(b_i) : r \in \{row(i,j) : 1 \leq j \leq |B_i|\}\} = \{\pi_{row(i,j)}(b_i) : 1 \leq j \leq |B_i|\}$. Due to the construction of row, this is B_i. It follows that $|\mathrm{MAP}_{\prod^\ell, b_i, R_i}|$ is in fact $|B_i|$, which by the definition of R_i, and the way A is built, is also $|R_i|$. As the given S is a filtering set, it follows that $B_i \subseteq \overline{N_G(a_i)}$. It also follows that $b_i \notin B_i$, that is $b_i \notin \mathrm{MAP}_{\prod^\ell, b_i, R_i}$. Thus, condition (2) of Definition 3 is met. It is left to prove condition (3) in Definition 3. Let $r \in R_i$, we want to prove that $\mathrm{MAP}_{\prod^\ell, b_i, R_i} \bigcup \{b_i\} \setminus \{\pi_r(b_i)\} \subseteq \overline{N_{G_{\pi_r}}(a_i)}$. That is, $B_i \bigcup \{b_i\} \setminus \{\pi_r(b_i)\} \subseteq \overline{N_{G_{\pi_r}}(a_i)}$. Which is equivalent to $B_i \bigcup \{b_i\} \setminus \{\pi_r(b_i)\} \subseteq \overline{\pi_r(N_G(a_i))}$ (We use the fact that $\pi_r(a_i) = a_i$, as this is how A was built). The way A is built makes sure that $\pi_r(B_i) = B_i \bigcup \{b_i\} \setminus \{\pi_r(b_i)\}$. So we need to show that $\pi_r(B_i) \subseteq \pi_r(N_G(a_i))$. This simply follows from $B_i \subseteq \overline{N_G(a_i)}$).

Let $(a_{i_1}, b_{i_1}, R_{i_1})$, $(a_{i_2}, b_{i_2}, R_{i_2})$ be two distinct triplets in S^* and let $(a_{i_1}, b_{i_1}, B_{i_1})$, $(a_{i_2}, b_{i_2}, B_{i_2})$ be the respective triplets in S. That is, $a_{i_1} \neq b_{i_1}$ and $a_{i_2} \neq b_{i_2}$. To prove that S^* is a $\prod_G^\ell$-multiplexing set, we need to prove that $a_{i_1} \notin \{b_{i_2}\} \bigcup \mathrm{MAP}_{\prod^\ell, b_{i_2}, R_{i_2}}$ and $(\{b_{i_1}\} \bigcup \mathrm{MAP}_{\prod^\ell, b_{i_1}, B_{i_1}}) \bigcap (\{b_{i_2}\} \bigcup \mathrm{MAP}_{\prod^\ell, b_{i_2}, R_{i_2}}) = \emptyset$. The conditions translate to $a_{i_1} \notin \{b_{i_2}\} \bigcup B_{i_2}$ and $\{b_{i_1}\} \bigcup B_{i_1}) \bigcap (\{b_{i_2}\} \bigcup B_{i_2}) = \emptyset$, which is true as S is a filtering set.

We can now apply Theorem 1, and get the desired result. □

The following example demonstrates the construction of a matrix A as in the above proof.

Example 2. Let $G = (V, E)$ be a restriction graph for 9 parties, where the set of edges $E = \{(1,2),(1,5),(7,8)\}$

Take the ordered set $S_1 = (1,2,\{3,4\}), S_2 = (1,5,\{6\}), S_3 = (7,8,\{9\})$. As $R(1) = 3, R(7) = 1$ and both are bound by $\ell - 1 = 3$. So the set is a 4-filtering set of size 3. The appropriate $row's$ are $row(1,1) = 2, row(1,2) = 3, row(2,1) = 4, row(3,1) = 2$. The corresponding matrix is

$$A = \begin{pmatrix} 1\,2\,3\,4\,5\,6\,7\,8\,9 \\ 1\,3\,?\,?\,?\,?\,?\,9\,? \\ 1\,4\,?\,?\,?\,?\,?\,?\,? \\ 1\,?\,?\,?\,6\,?\,?\,?\,? \end{pmatrix}$$

Example 3. We show that Theorem 2 implies Proposition 2. Let k be as in Proposition 2. Let $G = ([k], [k] \times [k] \setminus \{(i,i) : i \in [k]\} \setminus \{(k,4),\ldots,(k,k-1)\})$ be a restriction graph. I.e, G is a complete digraph, except that it doesn't have edges from the node k to any of the nodes in $[4, k-1]$. We describe a NOF_G protocol Q for f: party P_k sends P_1 the bit 1 if $x_1 = x_2$, and 0 otherwise. If P_2 got 1 and all other inputs are equal, it outputs 1, otherwise it outputs 0. For the $\frac{k-3}{2}$-size $(\frac{k-1}{2})$-filtering set $S = \{(k, 2, \{4, 6, \ldots k-1\})\}$, we apply Theorem 2. It follows that $D^{\mathrm{NOF}}(EQ_k^{\frac{k-1}{2}}) \leq \frac{k-1}{2} \cdot 1 + \frac{k-1}{2} - 1 \cdot (\frac{k-3}{2}) = 1 + \frac{k-1}{2}$.

4 Multiplexing NOF protocols via the Myopic Multi-Party Model

In this section, we show that given so-called MYOPIC protocols for ℓ permutations of their k parties, that satisfy certain combinatorial requirements, we can "multiplex" those protocols into one NOF protocol, obtaining some savings in communication complexity. We start with some definitions.

Definition 10. *The k-party one-way* NOF *model is identical to the* NOF *model, except that the communication pattern is restricted to the following "one-way" manner: the parties are organized according to some permutation π and each party, in its turn, sends a single message to the next party according to π (namely, party $P_{\pi(i)}$ sends a message to $P_{\pi(i+1)}$). The last party announces the output. In the* myopic *one-way model, there is an additional restriction that party $P_{\pi(i)}$ sees only the inputs of the pervious parties and the following party. That is, $P_{\pi(1)}, \ldots, x_{\pi(i-1)}$ and $x_{\pi(i+1)}$. For a permutation π over $[k]$, we denote MYOPIC_π as the myopic model with π being the order of parties.*

As an example of a myopic protocol, see the protocol in the proof of Proposition 1, where instead of P_{k-1} broadcasting its message, it sends its message to P_k and P_k announces the output.

The myopic model was first discussed in [16], where it was investigated with regards to the Pointer Jumping problem. The additional restriction, compared to the one-way model, was used to show better lower bounds on the communication complexity of

randomized protocols. It was further studied in [8], where better lower bounds were shown.

A MYOPIC_π protocol Q, can be thought of as a special case of a restricted-NOF protocol, for the graph G_π, where $G = ([k], \{(i,j) : (j < i) \vee (j = i+1)\})$. In addition, another restriction on Q, is that messages are sent in a "one-way" fashion, according to π. That is, if $|m^Q_{i\to j,t}| > 0$ for some input $\boldsymbol{x}$, then $\pi(t) = i, \pi(t+1) = j$. Therefore, the next result may be thought of as corollary of Theorem 2, except that, we do not require the protocol Q to be oblivious (nor do we require that other protocols which will be used, will have the same communication pattern as Q). Thus, we get a more general result.

We now describe, informally, which restrictions are required of two MYOPIC protocols Q_{π_1}, Q_{π_2} to be able to multiplex messages sent in them. Assume we have k parties and we want to multiplex messages sent by P_i. So, P_i needs to appear as the t^{th} party for some $t \in [k-1]$, in Q_{π_1}, Q_{π_2}. We also need the parties which receive the messages sent by P_i in those permutations to be different. That is, we need $|\{\pi_1(t+1), \pi_2(t+1)\}| = 2$. Next, in analogy to the requirements for multiplexing in the restricted model, the message P_i sends $P_{\pi_1(t+1)}$ in Q_{π_1}, should not depend on $x_{\pi_2(t+1)}$ (similarly for $P_{\pi_2(t+1)}$ in Q_{π_2} and $x_{\pi_1(t+1)}$). To accomplish this, we require that $P_{\pi_2}(t+1)$ does not appear in π_1 before the t^{th} place. Lastly, we need that messages sent to P_i in Q_{π_1} and Q_{π_2}, will not be multiplexed, so in the combined NOF protocol, these messages will appear plainly on the board, so $P_{\pi_1(t+1)}, P_{\pi_2(t+1)}$ can see them. These requirements can be extended to accommodate multiplexing ℓ protocols.

Since we don't assume that Q_{π_1}, Q_{π_2} were oblivious, the savings we have made in the communication complexity, are per input $\boldsymbol{x}$, and therefore the resulting Theorem has a different expression than the one in Theorem 1, for the bound on communication complexity of the direct-sum.

We formulate the necessary requirements for MYOPIC protocols to be multiplexed, in an analogy to Definition 4.

Definition 11. *Let $\prod^\ell = \pi_1, \ldots, \pi_\ell$ be a sequence of ℓ permutations over $[k]$. A triplet (i, s, U) is $\prod^\ell$-binding if $i, s \in [k]$ and $U \subseteq [\ell]$ such that:*

1. *For all $u \in U : \pi_u(i) = s$.*
2. *$|\pi_u(i+1) : u \in U| = |U|$.*
3. *$\{\pi_u(i') : u \in U, i' < i\} \bigcap \{\pi_u(i+1) : u \in U\} = \emptyset$.*

The definition guarantees a sequence of $|U|$ permutations, such that for each of the permutations, in step t, the party P_s is the one who sends a message. Also guaranteed, is that all the recipients of those messages, according to the $|U|$ permutations, are parties which have not appeared before P_s in any of the $|U|$ permutations.

Definition 12. *Let $\prod^\ell$ be a sequence of ℓ permutations over $[k]$. A $\prod^\ell$-repetitive set S is a set of $\prod^\ell$-binding triplets, such that for all distinct $(t_1, s_1, U_1), (t_2, s_2, U_2) \in S$:*

1. *$s_1 \notin \{\pi_u(t_2+1) : u \in U_2\}$.*
2. *if $(t_1, s_1) = (t_2, s_2)$ then $U_1 \bigcap U_2 = \emptyset$.*

Theorem 3. *Let $\prod^{\ell}$ be a sequence of ℓ permutations over $[k]$. Let f be a k-argument function and S a $\prod^{\ell}$-repetitive set. Let $Q_{\pi_1},\ldots,Q_{\pi_\ell}$ be ℓ protocols corresponding to the computation of f in MYOPIC_{π_i}. Also assume that messages sent during those protocols are prefix-free. Denote $Cost_{\boldsymbol{x},u,t}$ as the number of bits sent in the message from $\pi_u(t)$ to $\pi_u(t+1)$, in protocol Q_{π_u}, on input $\boldsymbol{x}$. For a $\prod^{\ell}$-binding triplet (t,s,U) denote $TCost_{\boldsymbol{x},U,t}=\sum_{u\in U}Cost_{\boldsymbol{x},u,t}$, meaning, the total number of bits sent from party P_s to parties in $\{\pi(u)(t+1):u\in U\}$, in the corresponding protocols. Also denote $MCost_{\boldsymbol{x},U,t}=\max_{u\in U}Cost_{\boldsymbol{x},u,t}$, as the maximal length of a message, from the aforementioned messages. Then,*

$$D^{\mathrm{NOF}}(f^{\ell})\leq\max_{\boldsymbol{x}}(\sum_{u=1\ldots l}\sum_{t=1\ldots\ell-1}Cost_{\boldsymbol{x},u,t}-\sum_{(a,s,U)\in S}(TCost_{\boldsymbol{x},U,a}-MCost_{\boldsymbol{x},U,a})).$$

Proof. We describe a protocol Q' for f^{ℓ} in the NOF model. We will use protocol Q_i to solve f on the i^{th} instance of f^{ℓ}. Protocol Q' will run all those protocols in parallel and will perform multiplexing when possible. In round t, for each instance i, the corresponding step of protocol Q_i is made on the i^{th} instance, only that instead of actually sending the message on a private channel, as in the myopic model, the corresponding party writes it on the board. To reduce the amount of communication, when performing round t, for every $(t,s,U)\in S$, instead of P_s sending $|U|$ messages to each $\pi_u(t+1)$ for $u\in U$, party P_s xors those $|U|$ messages together, and writes on the board the xored message. We argue that each of the original recipients can then restore the original message it was supposed to receive. We take $u_1\in U$, and show that party $\pi_{u_1}(t+1)$ can retrieve its message. We need to show that for each $u\in\{U\backslash u_1\}$, party $\pi_{u_1}(t+1)$ can compute the message sent at round t from P_s to $P_{\pi_u(t+1)}$ in Q_u. This message may depend, by the definition of MYOPIC_{π_u}, on inputs of $\{P_{\pi_u(t')}|t'\leq t+1,t'\neq t\}$ and on the message received by P_s in Q_u. By condition (3) of the definition of a $\prod^{\ell}$-repetitive set, this is disjoint to $\{P_{\pi_{u_1}(t+1)}\}$ (as $u,u_1\in U$). That is, they are known to $P_{\pi_{u_1}(t+1)}$ in the NOF model. The received message of P_s in Q_u is not multiplexed (as for all $(t,s,U)\neq(t_2,s_2,U_2)\in S$ it is promised that $s\notin\{\pi_u(t_2+1):u\in U_2\}$), and therefore written on the board. It follows that these messages can be computed by u_1 in the NOF model, and that it can also compute its original message, by xoring all of them with the xored valued on the board. Note that for the case when its original message was shorter then the other multiplexed messages, it still need to know how many bits of the xored result it needs to take. For that we use the fact that messages in the protocol are prefix-free.

That is, for every $(t,s,U)\in S$, we manage to send a single message only, instead of $|U|$, when the length of the message sent is the longest massage out of the possible $|U|$ (the shorter ones can be padded to fit the length of the longest one). □

Corollary 2. *Let $\prod^{\ell}$ be a sequence of ℓ permutations over $[k]$. Let f be a k-argument function and S a $\prod^{\ell}$-repetitive set. Let $Q_{\pi_1},\ldots,Q_{\pi_\ell}$ be ℓ oblivious protocols corresponding to the computation of f in MYOPIC_{π_i}, such that they all have the same communication pattern. Then,*

$$D^{\mathrm{NOF}}(f^{\ell})\leq\ell\cdot D^{\mathrm{MYOPIC}_{\pi_1}}(Q_1)+\ell-\sum_{(a,b,U)\in S}(|U|-1)\cdot W_{Q_1}(a,b).$$

Note that Corollary 2 can also be viewed as a a Corollary to Theorem 1, as for the case of oblivious protocols, Theorem 3 is a private case of Theorem 1.

References

1. A. Ada, A. Chattopadhyay, O. Fawzi, and P. Nguyen *The NOF Multiparty Communication Complexity of Composed Functions*. In Proceedings of the 39th International Colloquium on Automata, Languages, and Programming (ICALP), volume 7391 of Lecture Notes in Computer Science, pages 13-24. Springer, 2012.
2. A. Ambainis, H. Buhrman, W. Gasarch, B. Kalyanasundaramd, and L. Torenvliete. *The communication complexity of enumeration, elimination, and selection*. J. Comput. Syst. Sci., 63(2):148-185, 2001.
3. L. Babai, N. Nisan, and M. Szegedy. *Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs*. JCSS, 45, 1992.
4. B. Barak, M. Braverman, X. Chen, A. Rao. *Direct Sums in Randomized Communication Complexity*. Electronic Colloquium on Computational Complexity (ECCC) 16 - 44. 2009.
5. P. Beame, M. David, T. Pitassi, and P. Woelfel. *Separating deterministic from randomized multiparty communication complexity*. Theory of Computing, 6(1): pages 201225, 2010.
6. P. Beame, T. Pitassi, N. Segerlind, and A. Wigderson. *A strong direct product theorem for corruption and the multiparty communication complexity of Disjointness*. Computational Complexity, 15(4):391432, 2006.
7. A.Beimel, S. Ben Daniel, E. Kushilevitz, and E. Weinreb. *Choosing, agreeing, and eliminating in communication complexity*. In 37th ICALP, pages 451-462, 2010.
8. A. Chakrabarti. *Lower bounds for multi-player pointer jumping*. In IEEE Conf. Computational Complexity (CCC), pages 3345, 2007.
9. A. Chattopadhyay and A. Ada. *Multiparty communication complexity of disjointness*. Technical Report TR-08-002, ECCC, 2008.
10. A. K. Chandra, M. L. Furst, R. J.Lipton, *Multi-party Protocols*. Proc. 15th STOC, pages 94-99, 1983.
11. M. David, T. Pitassi, and E. Viola. *Improved separations between nondeterministic and randomized multiparty communication*. In APPROX-RANDOM, volume 5171 of Lecture Notes in Computer Science, pages 371384. Springer, 2008.
12. J. S. Ford *Lower Bound Methods for Multiparty Communication Complexity*. PhD Thesis, The University of Texas at Austin. 121 p.
13. J. Draisma, E. Kushilevitz, and E. Weinreb. *Partition arguments in multiparty communication complexity.* Theor. Comput. Sci., Vol 412(24), pp. 2611-2622, 2011. Early version in ICALP 2009.
14. T. Feder, E. Kushilevitz, M. Naor, and N. Nisan. *Amortized communication complexity*. SIAM Journal on Computing, 24(4):736-750, 1995. Early version in FOCS'91.
15. V. Grolmusz *On multi-party communication complexity of random functions*. MPI-I-93-162. December Max Planck Institut fuer Informatik. 1993.
16. A. Gronemeier. *NOF-multiparty information complexity bounds for pointer jumping*. In Proc. 31st International Symposium on Mathematical Foundations of Computer Science, 2006.
17. R. Jain. *New strong direct product results in communication complexity*. Electronic Colloquium on Computational Complexity (ECCC) 18 - 24. 2011.
18. R. Jain, A. Pereszlnyi, P. Yao. *A direct product theorem for bounded-round public-coin randomized communication complexity*. CoRR abs/1201 - 1666. 2012.

19. M. Karchmer, R. Raz, and A. Wigderson. *Super-logarithmic depth lower bounds via the direct sum in communication complexity*. Computational Complexity, 5(3/4):191 - 204. 1995. Early version in CCC 1991.
20. M. Karchmer, E. Kushilevitz, and N. Nisan. *Fractional covers and communication complexity*. In Proceedings, Structure in Complexity Theory, 7th Annual Conference, pages 262-274. 1992. IEEE.
21. E. Kushilevitz and N. Nisan. "Communication Complexity", Cambridge University Press. 1997.
22. T. Lee and A. Shraibman. *Disjointness is hard in the multi-party number in the forehead model*. In Electronic Colloquium on Computational Complexity, number TR08-003, 2008.
23. J. M. Phillips, E. Verbin, and Qin Zhang. *Lower bounds for Number-in-Hand multiparty communication complexity, made easy*. Proc. of SODA 2012, pp. 486-501.
24. A. Sherstov. *The multiparty communication complexity of set disjointness*. Proc. of 44th STOC, 2012.
25. A. C. Yao. *Some Complexity Questions Related to Distributive Computing*. Proc. 11th STOC, pages. 209-213. 1979.